

Efficient routing of multiple vehicles with no communications

Alessandro Arsie and Emilio Frazzoli *

Laboratory for Information and Decision Systems,
Massachusetts Institute of Technology, Cambridge, MA
arsie@ucla.edu, frazzoli@mit.edu

February 2, 2008

Abstract

In this paper we consider a class of dynamic vehicle routing problems, in which a number of mobile agents in the plane must visit target points generated over time by a stochastic process. It is desired to design motion coordination strategies in order to minimize the expected time between the appearance of a target point and the time it is visited by one of the agents. We propose control strategies that, while making minimal or no assumptions on communications between agents, provide the same level of steady-state performance achieved by the best known decentralized strategies. In other words, we demonstrate that inter-agent communication does not improve the efficiency of such systems, but merely affects the rate of convergence to the steady state. Furthermore, the proposed strategies do not rely on the knowledge of the details of the underlying stochastic process. Finally, we show that our proposed strategies provide an efficient, pure Nash equilibrium in a game theoretic formulation of the problem, in which each agent's objective is to maximize the number of targets it visits. Simulation results are presented and discussed.

1 Introduction

A very active research area today addresses coordination of several mobile agents: groups of autonomous robots and large-scale mobile networks are being considered for a broad class of applications, ranging from environmental monitoring, to search and rescue operations, and national security.

An area of particular interest is concerned with the generation of efficient co-operative strategies for several mobile agents to move through a certain number of given target points, possibly avoiding obstacles or threats [1–5]. Trajectory

*Research supported by National Science Foundation, and Air Force Office of Scientific Research, CCR-0325716, and F49620-02-1-0325, respectively.

efficiency in these cases is understood in terms of cost for the agents: in other words, efficient trajectories minimize the total path length, the time needed to complete the task, or the fuel/energy expenditure. A related problem has been investigated as the Weapon-Target Assignment (WTA) problem, in which mobile agents are allowed to team up in order to enhance the probability of a favorable outcome in a target engagement [6, 7]. In this setup, targets locations are known and an assignment strategy is sought that maximizes the global success rate. In a biological setting, the closest parallel to many of these problems is the development of foraging strategies, and of territorial vs. gregarious behaviors [8], in which individuals choose to identify and possibly defend a hunting ground.

In this paper we consider a class of cooperative motion coordination problems, to which we can refer as *dynamic vehicle routing*, in which service requests are not known a priori, but are dynamically generated over time by a stochastic process in a geographic region of interest. Each service request is associated to a target point in the plane, and is fulfilled when one of a team of mobile agents visits that point. For example, service requests can be thought of as threats to be investigated in a surveillance application, events to be measured in an environmental monitoring scenario, and as information packets to be picked up and delivered to a user in a wireless sensor network. It is desired to design a control strategy for the mobile agents that provably minimizes the expected waiting time between the issuance of a service request and its fulfillment. In other words, our focus is on the quality of service as perceived by the “end user,” rather than, for example, fuel economies achieved by the mobile agents. Similar problems were also considered in [9, 10], and decentralized strategies were presented in [11]. This problem has connections to the Persistent Area Denial (PAD) and area coverage problems discussed, e.g., in [3, 12–14].

A common theme in cooperative control is the investigation of the effects of different communication and information sharing protocols on the system performance. Clearly, the ability to access more information at each single agent can not decrease the performance level; hence, it is commonly believed that by providing better communication among agents will improve the system’s performance. In this paper, we prove that there are certain dynamic vehicle routing problems which can, in fact, be solved (almost) optimally without any explicit communication between agents; in other words, the no-communication constraint in such cases is not binding, and does not limit the steady-state performance. The main contribution of this paper is the introduction of a motion coordination strategy that does not require any explicit communication between agents, while achieving provably optimal performance in certain conditions.

The paper is structured as follows: in Section 2 we set up and formulate the problem we investigate in the paper. In Section 3 we introduce the proposed solution algorithms, and discuss their characteristics. Section 4 is the technical core of the paper, in which we prove the convergence of the performance provided by the proposed algorithms to a critical point (either a local minimum or a saddle point) of the global performance function. Moreover, we show that any optimal configuration corresponds to a class of tessellations of

the plane that we call Median Voronoi Tessellations. Section 5 is devoted to a game-theoretic interpretation of our result in which the agents are modeled as rational autonomous decision makers trying to maximize their own utility function. We prove that, following the policy prescribed by our algorithm, the agents reach an efficient pure Nash equilibrium which can be anyway suboptimal with respect to the global utility function (in this case the expected time of service). In Section 6 we present some numerical results, while Section 7 is dedicated to final remarks and further extensions of this line of research.

2 Problem Formulation

Let $\Omega \subset \mathbb{R}^2$ be a convex domain on the plane, with non-empty interior; we will refer to Ω as the *workspace*. A stochastic process generates *service requests* over time, which are associated to points in Ω ; these points are also called *targets*. The process generating service requests is modeled as a spatio-temporal Poisson point process, with temporal intensity $\lambda > 0$, and an absolutely continuous spatial distribution described by the density function $\varphi : \Omega \rightarrow \mathbb{R}_+$, with bounded and convex support within Ω (i.e., $\varphi(q) > 0 \Leftrightarrow q \in \mathcal{Q} \subseteq \Omega$, with \mathcal{Q} bounded and convex). The spatial density function φ is normalized in such a way that $\int_{\Omega} \varphi(q) dq = 1$. Both λ and φ are not necessarily known.

A spatio-temporal Poisson point process is a collection of functions $\{\mathcal{P} : \overline{\mathbb{R}}_+ \rightarrow 2^{\Omega}\}$ such that, for any $t > 0$, $\mathcal{P}(t)$ is a random collection of points in Ω , representing the service requests generated in the time interval $[0, t)$, and such that

- The total numbers of events generated in two disjoint time-space regions are *independent* random variables;
- The total number of events occurring in an interval $[s, s+t)$ in a measurable set $\mathcal{S} \subseteq \Omega$ satisfies

$$\Pr[\text{card}((\mathcal{P}(s+t) - \mathcal{P}(s)) \cap \mathcal{S}) = k] = \frac{\exp(-\lambda t \cdot \varphi(\mathcal{S}))(\lambda t \cdot \varphi(\mathcal{S}))^k}{k!},$$

where this must hold for any k in \mathbb{N} and where $\varphi(\mathcal{S})$ is a shorthand for $\int_{\mathcal{S}} \varphi(q) dq$.

Each particular function \mathcal{P} is a realization, or trajectory, of the Poisson point process. A consequence of the properties defining Poisson processes is that the *expected* number of targets generated in a measurable region $\mathcal{S} \subseteq \Omega$ during a time interval of length Δt is given by:

$$\mathbb{E}[\text{card}((\mathcal{P}(t + \Delta t) - \mathcal{P}(t)) \cap \mathcal{S})] = \lambda \Delta t \cdot \varphi(\mathcal{S}).$$

Without loss of generality, we will identify service requests with targets points, and label them in order of generation; in other words, given two targets $e_i, e_j \in \mathcal{P}(t)$, with $i < j$, the service request associated with these target have been

issued at times $t_i \leq t_j \leq t$ (since events are almost never generated concurrently, the inequalities are in fact strict almost surely).

A service request is fulfilled when one of m mobile agents, modeled as point masses, moves to the target point associated with it; m is a possibly large, but finite number. Let $p(t) = (p_1(t), p_2(t), \dots, p_m(t)) \in \Omega^m$ be a vector describing the positions of the agents at time t . (We will tacitly use a similar notation throughout the paper). The agents are free to move, with bounded speed, within the workspace Ω ; without loss of generality, we will assume that the maximum speed is unitary. In other words, the dynamics of the agents are described by differential equations of the form

$$\frac{d p_i(t)}{dt} = u_i(t), \quad \text{with } \|u_i(t)\| \leq 1, \quad \forall t \geq 0, i \in \{1, \dots, m\}. \quad (1)$$

The agents are identical, and have unlimited range and target-servicing capability.

Let $\mathcal{B}_i(t) \subset \Omega$ indicate the set of targets serviced by the i -th agent up to time t . (By convention, $\mathcal{B}_i(0) = \emptyset$, $i = 1, \dots, m$). We will assume that $\mathcal{B}_i \cap \mathcal{B}_j = \emptyset$ if $i \neq j$, i.e., that service requests are fulfilled by at most one agent. (In the unlikely event that two or more agents visit a target at the same time, the target is arbitrarily assigned to one of them).

Let $\mathcal{D} : t \rightarrow 2^\Omega$ indicate (a realization of) the stochastic process obtained combining the service request generation process \mathcal{P} and the removal process caused by the agents servicing outstanding requests; in other words,

$$\mathcal{D}(t) = \mathcal{D}(t) \cup \mathcal{B}_1(t) \cup \dots \cup \mathcal{B}_m(t), \quad \mathcal{D}(t) \cap \mathcal{B}_i(t) = \emptyset, \quad \forall i \in \{1, \dots, m\}.$$

The random set $\mathcal{D}(t) \subset \Omega$ represents the *demand*, i.e., the service requests outstanding at time t ; let $n(t) = \text{card}(\mathcal{D}(t))$.

Our objective in this paper will be the design of motion coordination strategies that allow the mobile agents to fulfill service requests efficiently (we will make this more precise in the following). In particular, in this paper we will concentrate on motion coordination strategies of the following two forms:

$$\pi_i : (p_i, \mathcal{B}_i, \mathcal{D}) \mapsto u_i, \quad i \in \{1, \dots, m\}, \quad (2)$$

and

$$\pi_i : (p_1, \dots, p_m, \mathcal{B}_i, \mathcal{D}) \mapsto u_i, \quad i \in \{1, \dots, m\}. \quad (3)$$

An agent executing a control policy of the form (2) relies on the knowledge of its own current position, on a record of targets it has previously visited, and on the current demand. In other words, such control policies do not need any explicit information exchange between agents; as such, we will refer to them as *no communication* (nc) policies. Such policies are trivially decentralized.

On the other hand, an agent executing a control policy of the form (3) can sense the current position of other agents, but still has information only on the targets itself visited in the past (i.e., does not know what, if any, targets have been visited by other agents). We call these *sensor-based* (sb) policies, to signify

the fact that only factual information is exchanged between agents—as opposed to information related to intent and past history. Note that both families of coordination policies rely, in principle, on the knowledge of the locations of all outstanding targets. (However, as we will see in the following, only local target sensing will be necessary in practice).

A policy $\pi = (\pi_1, \pi_2, \dots, \pi_m)$ is said to be *stabilizing* if, under its effect, the expected number of outstanding targets does not diverge over time, i.e., if

$$\bar{n}_\pi = \lim_{t \rightarrow \infty} \mathbb{E}[n(t) | \dot{p}_i(t) = \pi_i(p(t), \mathcal{B}_i(t), \mathcal{D}(t)), i \in \{1, \dots, m\}] < \infty. \quad (4)$$

Intuitively, a policy is stabilizing if the mobile agents are able to visit targets at a rate that is—on average—at least as fast as the rate at which new service requests are generated.

Let T_j be the time elapsed between the issuance of the j -th service request, and the time it is fulfilled. If the system is stable, then the following balance equation (also known as Little’s formula [15]) holds:

$$\bar{n}_\pi = \lambda \bar{T}_\pi, \quad (5)$$

where $\bar{T}_\pi := \lim_{j \rightarrow \infty} \mathbb{E}[T_j]$ is the system time under policy π , i.e., the expected time a service request must wait before being fulfilled, given that the mobile agents follow the strategy defined by π . Note that the system time \bar{T}_π can be thought of as a measure of the quality of service, as perceived by the “user” issuing the service requests.

At this point we can finally state our problem: we wish to devise a policy that is (i) stabilizing, and (ii) yields a quality of service (i.e., system time) achieving, or approximating, the theoretical optimal performance given by

$$\bar{T}_{\text{opt}} = \inf_{\pi \text{ stabilizing}} \bar{T}_\pi \quad (6)$$

Centralized and decentralized strategies are known that optimize or approximate (6) in a variety of cases of interest [10, 11, 16, 17]. However, all such strategies rely either on a central authority with the ability to communicate to all agents, or on the exchange of certain information about each agent’s strategy with other neighboring agents. In addition, these policies require the knowledge of the spatial distribution φ ; decentralized versions of these implement versions of Lloyd’s algorithm for vector quantization [18].

In the remainder of this paper, we will investigate how the additional constraints posed on the exchange of information between agents by the models (2) and (3) impact the achievable performance and quality of service. Remarkably, the policies we will present do not rely on the knowledge of the spatial distribution φ , and are a generalized version of MacQueen’s clustering algorithm [19].

3 Control policy description

In this section, we introduce two control policies of the forms, respectively, (2) and (3). An illustration of the two policies is given in Figure 1.




Figure 1: An illustration of the two control policies proposed in Section 3. While no targets are outstanding, vehicles wait at the point that minimizes the average distance to targets they have visited in the past; such points are depicted as squares, while targets are circles and vehicles triangles. In the no-communication policy, at the appearance of a new target, all vehicles pursue it (left). In the sensor-based policy, only the vehicle that is closest to the target will pursue it (right).

3.1 A control policy requiring no explicit communication

Let us begin with an informal description of a policy π_{nc} requiring no explicit information exchange between agents. At any given time t , each agent computes its own control input according to the following rule:

1. If $\mathcal{D}(t)$ is not empty, move towards the nearest outstanding target.
2. If $\mathcal{D}(t)$ is empty, move towards the point minimizing the average distance to targets *served in the past* by each agent. If there is no unique minimizer, then move to the nearest one.

In other words, we set

$$\pi_{\text{nc}}(p_i(t), \mathcal{B}_i(t), \mathcal{D}(t)) = \text{vers}(F_{\text{nc}}(p_i(t), \mathcal{B}_i(t), \mathcal{D}(t)) - p_i(t)), \quad (7)$$

where

$$F_{\text{nc}}(p_i, \mathcal{B}_i, \mathcal{D}) = \begin{cases} \arg \min_{q \in \mathcal{D}} \|p_i - q\|, & \text{if } \mathcal{D} \neq \emptyset, \\ \arg \min_{q \in \Omega} \sum_{e \in \mathcal{B}_i} \|e - q\|, & \text{otherwise,} \end{cases} \quad (8)$$

$\|\cdot\|$ is the Euclidean norm, and

$$\text{vers}(v) = \begin{cases} v/\|v\|, & \text{if } v \neq 0, \\ 0 & \text{otherwise.} \end{cases}$$

The convex function $W : q \mapsto \sum_{e \in \mathcal{B}} \|q - e\|$, often called the (discrete) Weber function in the facility location literature [20, 21] (modulo normalization by

$\text{card}(\mathcal{B}))$, is not strictly convex only when the point set \mathcal{B} is empty—in which case we set $W(\cdot) = 0$ by convention—or contains an even number of collinear points. In such cases, the minimizer nearest to p_i in (8) is chosen. We will call the point $p_i^*(t) = F_{\text{nc}}(\cdot, \mathcal{B}_i(t), \emptyset)$ the *reference point* for the i -th agent at time t .

In the π_{nc} policy, whenever one or more service requests are outstanding, all agents will be pursuing a target; in particular, when only one service request is outstanding, all agents will move towards it. When the demand queue is empty, agents will either (i) stop at the current location, if they have visited no targets yet, or (ii) move to their reference point, as determined by the set of targets previously visited.

3.2 A sensor-based control policy

The control strategy in the previous section can be modified to include information on the current position of other agents, if available (e.g., through on-board sensors). In order to present the new policy, indicate with $\mathcal{V}(p) = \{\mathcal{V}_1(p), \mathcal{V}_2(p), \dots, \mathcal{V}_m(p)\}$ the Voronoi partition of the workspace Ω , defined as:

$$\mathcal{V}_i(p) = \{q \in \Omega : \|q - p_i\| \leq \|q - p_j\|, \forall j = 1 \dots m\}. \quad (9)$$

As long as an agent has never visited any target, i.e., as long as $\mathcal{B}_i(t) = \emptyset$, it executes the π_{nc} policy. Once an agent has visited at least one target, it computes its own control input according to the following rule:

1. If $\mathcal{D}(t) \cap \mathcal{V}_i(t)$ are not empty, move towards the nearest outstanding target in the agent's own Voronoi region.
2. If $\mathcal{D}(t) \cap \mathcal{V}_i(t)$ is empty, move towards the point minimizing the average distance to targets in $\mathcal{B}_i(t)$. If there is no unique minimizer, then move to the nearest one.

In other words, we set

$$\pi_{\text{sb}}(p(t), \mathcal{B}_i(t), \mathcal{D}(t)) = \text{vers}(F_{\text{sb}}(p(t), \mathcal{B}_i(t), \mathcal{D}(t)) - p_i(t)), \quad (10)$$

where

$$F_{\text{sb}}(p, \mathcal{B}_i, \mathcal{D}) = \begin{cases} \arg \min_{q \in \mathcal{D}} \|p_i - q\|, & \text{if } \mathcal{D} \cap \mathcal{V}_i \neq \emptyset, \text{ and } \mathcal{B}_i = \emptyset \\ \arg \min_{q \in \mathcal{D} \cap \mathcal{V}_i(p)} \|p_i - q\|, & \text{if } \mathcal{D} \cap \mathcal{V}_i \neq \emptyset, \text{ and } \mathcal{B}_i \neq \emptyset \\ \arg \min_{q \in \Omega} \sum_{e \in \mathcal{B}_i} \|e - q\|, & \text{otherwise.} \end{cases} \quad (11)$$

In the π_{sb} policy, at most one agent will be pursuing a given target, at any time after an initial transient that terminates when all agents have visited at least one target each. The agents' behavior when no outstanding targets are available in their Voronoi region is similar to that determined by the π_{nc} policy previously discussed, i.e., they move to their reference point, determined by previously visited targets.

Remark 1 While we introduced Voronoi partitions in the definition of the control policy, the explicit computation of each agent’s Voronoi region is not necessary. In fact, each agent only needs to check whether it is the closest agent to a given target or not. In order to check whether a target point q is in the Voronoi region of the i -th agent, it is necessary to know the current position only of agents within a circle or radius $\|p_i - q\|$ centered at q (see Figure 2). For example, if such circle is empty, then q is certainly in \mathcal{V}_i ; if the circle is not empty, distances of the agents within it to the target must be compared. This provides a degree of spatial decentralization—with respect to other agents—that is even stronger than that provided by restricting communications to agents sharing a boundary in a Voronoi partition (i.e., neighboring agents in the Delaunay graph, dual to the partition (9)).




Figure 2: Implicit computation of Voronoi regions: Even though target e_1 is the nearest target to p_1 , it is not in the Voronoi region of the 1st agent. In fact, the circle of radius $\|e_1 - p_1\|$ centered at e_1 contains p_2 , and the 2nd agent is closer to e_1 . However, the circle of radius $\|e_2 - p_1\|$ centered at e_2 does not contain any other agent, ensuring that e_2 is in the Voronoi region generated by p_1 .

Remark 2 The sensor-based policy is more efficient than the no-communication policy in terms of the length of the path traveled by each agent, since there is no duplication of effort as several agents pursue the same target. However, in terms of “quality of service,” we will show that there is no difference between the two policies, for low target generation rates. Numerical results show that the sensor-based policy is more efficient in a broader range of target generation rates, and in fact provides almost optimal performance both in light and heavy load conditions.

4 Performance analysis in light load

In this section we analyze the performance of the control policies proposed in the previous section. In particular, we concentrate our investigation on the light load case, in which the target generation rate is very small, i.e., as $\lambda \rightarrow 0^+$. This will allow us to prove analytically certain interesting and perhaps surprising

characteristics of the proposed policies. The performance analysis in the general case is more difficult; we will discuss the results of numerical investigation in Section 6, but no analytical results are available at this time.

4.1 Overview of the system behavior in the light load regime

Before starting a formal analysis, let us summarize the key characteristics of the agents' behavior in light load, i.e., for small values of λ .

1. At the initial time the m agents are assumed to be deployed in general position in Ω , and the demand queue is empty, $\mathcal{D}(0) = \emptyset$.
2. The agents do not move until the first service request appears. At that time, if the policy π_{nc} is used, all agents will start moving towards the first target. If the sensor-based policy π_{sb} is used, only the closest agent will move towards the target.
3. As soon as one agent reaches the target, all agents start moving towards their current reference point, and the process continues.

For small λ , with high probability (i.e., with probability approaching 1 as $\lambda \rightarrow 0$) at most one service request is outstanding at any given time. In other words, new service requests are generated so rarely that most of the time agents will be able to reach a target and return to their reference point before a new service request is issued.

Consider the j -th service request, generated at time t_j . Assuming that at t_j all agents are at their reference position, the expected system time T_j can be computed as

$$T_j = \int_{\Omega} \min_{i=1,\dots,m} \|p_i^*(t_j) - q\| \varphi(q) dq.$$

Assume for now that the sequences $\{p_i^*(t_j) : j \in \mathbb{N}\}$ converge, and let

$$\lim_{j \rightarrow \infty} p_i^*(t_j) = \hat{p}_i^*.$$

Note that \hat{p}_i^* is a random variable, the value of which depends in general on the particular realization of the target generation process. If all service requests are generated with the agents at their reference position, the average service time (for small λ) can be evaluated as

$$\overline{T}_{\text{nc}} = \overline{T}_{\text{sb}} = \int_{\Omega} \min_{i=1,\dots,m} \|\hat{p}_i^* - q\| \varphi(q) dq = \sum_{i=1}^m \int_{\mathcal{V}_i(p^*)} \|\hat{p}_i^* - q\| \varphi(q) dq. \quad (12)$$

Since the system time depends on the random variable $\hat{p}^* = (\hat{p}_1^*, \dots, \hat{p}_m^*)$, it is itself a random variable. The function appearing on the right hand side of the above equation, relating the system time to the asymptotic location of reference points, is called the continuous multi-median function [20]. This function admits

a global minimum (in general not unique) for all non-singular density functions φ , and in fact it is known [10] that the optimal performance in terms of system time is given by

$$\bar{T}_{\text{opt}} = \min_{p \in \Omega^m} \sum_{i=1}^m \int_{\mathcal{V}_i(p)} \|p_i - q\| \varphi(q) dq. \quad (13)$$

In the following, we will investigate the convergence of the reference points as new targets are generated, in order to draw conclusions about the average system time \bar{T} in light load. In particular, we will prove not only that the reference points converge with high probability (as $\lambda \rightarrow 0$) to a local critical point (more precisely, either local minima or saddle points) for the average system time, but also that the limiting reference points \hat{p}^* are *generalized medians* of their respective Voronoi regions, where

Definition 3 (Generalized median) *The generalized median of a set $\mathcal{S} \subset \mathbb{R}^n$ with respect to a density function $\varphi : \mathcal{S} \rightarrow \mathbb{R}_+$ is defined as*

$$\bar{p} := \arg \min_{p \in \mathbb{R}^n} \int_{\mathcal{S}} \|p - q\| \varphi(q) dq.$$

We call the resulting Voronoi tessellation *Median Voronoi Tessellation* (MVT for short), in analogy with what is done with Centroidal Voronoi Tessellations. A formal definition is as follows:

Definition 4 (Median Voronoi Tessellation) *A Voronoi tessellation $\mathcal{V}(p) = \{\mathcal{V}_1(p), \dots, \mathcal{V}_m(p)\}$ of a set $\mathcal{S} \subset \mathbb{R}^n$ is said a Median Voronoi Tessellation of \mathcal{S} with respect to the density function φ if the ordered set of generators p is equal to the ordered set of generalized medians of the sets in $\mathcal{V}(p)$ with respect to φ , i.e., if*

$$p_i = \arg \min_{s \in \mathbb{R}^n} \int_{\mathcal{V}_i(p)} \|s - q\| \varphi(q) dq, \quad \forall i \in \{1, \dots, m\}.$$

Since the proof builds on a number of intermediate results, we provide an outline of the argument as a convenience to the reader.

1. First we prove that the reference point of any agents that visits an unbounded number of targets over time converges almost surely.
2. Second, we prove that, if $m \geq 1$ agents visit an unbounded number of targets over time, their reference points will converge to the generators of a MVT almost surely, as long as agents are able to return to their reference point infinitely often.
3. Third, we prove that all agents will visit an unbounded number of targets (this corresponds to a property of distributed algorithms that is often called *fairness* in computer science).
4. Finally, we prove that agents are able to return to their reference point infinitely often with high probability as $\lambda \rightarrow 0^+$.

Combining these steps, together with (6), will allow us to state that the reference points converge to a local critical point of the system time, with high probability as $\lambda \rightarrow 0^+$.

4.2 Convergence of reference points

Let us consider an agent i , such that

$$\lim_{t \rightarrow \infty} \text{card}(\mathcal{B}_i(t)) = \infty,$$

i.e., an agent that services an unbounded number of requests over time. Since the number of agents m is finite, and the expected number of targets generated over a time interval $[0, t)$ is proportional to t , at least one such agent will always exist. In the remainder of this section, we will drop the subscript i , since we will consider only this agent, effectively ignoring all others for the time being.

For any finite t , the set $\mathcal{B}(t)$ will contain a finite number of points. Assuming that $\mathcal{B}(t)$ contains at least three non-collinear points, the discrete Weber function $p \mapsto \sum_{q \in \mathcal{B}(t)} \|p - q\|$ is strictly convex, and has a unique optimizer $p^*(t) = \arg \min_{p \in \Omega} \sum_{q \in \mathcal{B}(t)} \|p - q\|$. The optimal point $p^*(t)$ is called the Fermat-Torricelli (FT) point—or the Weber point in the location optimization literature—associated with the set $\mathcal{B}(t)$; see [21–23] for a historical review of the problem and for solution algorithms.

It is known that the FT point is unique and algebraic for any set of non-collinear points. While there are no general analytic solutions for the location of the FT point associated to more than 4 points, numerical solutions can be easily constructed relying on the convexity of the Weber function, and on the fact that it is differentiable for all points not in \mathcal{B} . Polynomial-time approximation algorithms are also available (see, e.g., [24, 25]). Remarkably, a simple mechanical device can be constructed to solve the problem, based on the so-called Varignon frame, as follows. Holes are drilled on a horizontal board, at locations corresponding to the points in \mathcal{B} . A string attached to a unit mass is passed through each of these holes, and all strings are tied together at one end. The point reached by the knot at equilibrium is a FT point for \mathcal{B} .

Some useful properties of FT points are summarized below. If there is a $q_0 \in \mathcal{B}$ is such that

$$\left\| \sum_{q \in \mathcal{B} \setminus q_0} \text{vers}(q_0 - q) \right\| \leq 1 \quad (14)$$

then $p^* = q_0$ is a FT point for \mathcal{B} . If no point in \mathcal{B} satisfies such condition, then the FT point p^* can be found as a solution of the following equation:

$$\sum_{q \in \mathcal{B}} \text{vers}(p^* - q) = 0. \quad (15)$$

In other words, p^* is simply the point in the plane at which the sum of unit vectors starting from it and directed to each of the points in \mathcal{B} is equal to zero;

this point is unique if \mathcal{B} contains non-collinear points. Clearly, the FT point is in the convex hull of \mathcal{B} .

Note that the points in $\mathcal{B}(t)$ are randomly sampled from an unknown absolutely continuous distribution, described by a spatial density function $\tilde{\varphi}$ —which is not necessarily the same as φ , and in general is time-varying, depending on the past actions of all agents in the system. Even though $\tilde{\varphi}$ is not known, it can be expressed as

$$\varphi(q, t) = \begin{cases} \frac{\varphi(q)}{\int_{\mathcal{I}(t)} \varphi(q) dq} & \text{if } q \in \mathcal{I}(t) \\ 0 & \text{otherwise,} \end{cases}$$

for some convex set $\mathcal{I}(t)$ containing $p(t)$. (In practical terms, such set will be the Voronoi region generated by $p(t)$).

The function $t \mapsto p^*(t)$ is piecewise constant, i.e., it changes value at the discrete time instants $\{t_j : j \in \mathbb{N}\}$ at which the agent visits new targets. As a consequence, we can concentrate on the sequence $\{p^*(t_j) : j \in \mathbb{N}\}$, and study its convergence.

Definition 5 For any $t > 0$, let the solution set $C(t)$ be defined as

$$C(t) := \left\{ p \in \Omega : \left\| \sum_{q \in \mathcal{B}(t)} \text{vers}(p - q) \right\| \leq 1 \right\}.$$




Figure 3: Example of a Fermat-Torricelli point (star) and solution set corresponding to five target points (circles). Upon the addition of an arbitrarily chosen sixth target point, the Fermat-Torricelli is guaranteed to remain within the region bounded by the curve.

An example of such set is shown in Figure 3. The reason for introducing such solution sets is that they have quite remarkable properties as shown by the following

Proposition 6 For any $j \in \mathbb{N}$, $p^*(t_{j+1}) \in C(t_j)$. More specifically, if $e_{j+1} \notin C(t_j)$ (i.e., the target point associated to the j -th service request is outside the solution set) then the FT point $p^*(t_{j+1})$ is on the boundary of C_i . If $e_{j+1} \in C(t_j)$, then $p^*(t_{j+1}) = e_{j+1}$.

Proof: If e_{j+1} lies outside $C(t_j)$, we search for $p^*(t_{j+1})$ as the solution of the equation

$$\sum_{q \in \mathcal{B}(t_j)} \text{vers}(p - q) + \text{vers}(p - e_{j+1}) = 0,$$

from which it turns out immediately

$$\left\| \sum_{q \in \mathcal{B}(t_j)} \text{vers}(p - q) \right\| = \|\text{vers}(p - e_{j+1})\| = 1,$$

thus $p^*(t_{j+1}) \in \partial C(t_j)$. Notice that it is not true in general that the solution $p^*(t_{j+1})$ will lie on the line connecting $p^*(t_j)$ with the new target e_{j+1} . In the other case, if e_{j+1} lies in $C(t_j)$, then it satisfies condition (14), and is the new FT point. ■

Now, in order to prove that the $\{p^*(t_j)\}_{j \in \mathbb{N}}$ converges to a point \hat{p}^* , we will prove that the diameter of the solution set $C(t_j)$ vanishes almost surely as j tends to infinity. First we need the following result.

Proposition 7 If $\mathcal{Q} = \text{Supp}(\varphi)$ is convex with non-empty interior, then $p^*(t) \in \text{int}(\mathcal{Q})$ almost surely, for all t such that $\mathcal{B}(t) \neq \emptyset$.

Proof: For any non-empty target set $\mathcal{B}(t)$, the FT point lies within the convex hull of $\mathcal{B}(t)$. All points in the set $\mathcal{B}(t)$ are contained within the interior of \mathcal{Q} with probability one, since the boundary of \mathcal{Q} is a set of measure zero. Since $\text{int}(\mathcal{Q})$ is convex, and $\mathcal{B}(t) \subset \text{int}(\mathcal{Q})$ almost surely, $p^*(t) \in \text{co}(\mathcal{B}(t)) \subset \text{int}(\mathcal{Q})$, almost surely. ■

Proposition 8 If the support of φ is convex and bounded,

$$\lim_{j \rightarrow \infty} \text{diam}(C(t_j)) = 0, \quad \text{a.s.}$$

Proof: Consider a generic point $p \in C(t_j)$, and let $\delta = p - p^*(t_j)$,

$$\alpha_q = \arccos[\text{vers}(p - p^*(t_j)) \cdot \text{vers}(q - p^*(t_j))]$$

,

$$\alpha'_q = \arccos[\text{vers}(p - p^*(t_j)) \cdot \text{vers}(q - p)],$$

see Figure 4.




Figure 4: Geometric constructions in the proof of Proposition 8.

Since $p \in C(t_j)$, the magnitude of the sum of unit vectors $\sum_{q \in \mathcal{B}(t_j)} \text{vers}(p-q)$ is no more than one, and the following inequality is true:

$$\begin{aligned} & \left| \left(\sum_{q \in \mathcal{B}(t_j)} \text{vers}(p-q) \right) \cdot \text{vers}(p-p^*(t_j)) \right| \\ &= \left| \sum_{q \in \mathcal{B}(t_j)} (\text{vers}(p-q) \cdot \text{vers}(p-p^*(t_j))) \right| \\ &= \left| \sum_{q \in \mathcal{B}(t_j)} \cos(\alpha'_q) \right| \leq 1. \quad (16) \end{aligned}$$

Using elementary planar geometry, we obtain that

$$\alpha'_q - \alpha_q \geq \sin(\alpha'_q - \alpha_q) = \frac{\delta \sin(\alpha_q)}{\|q-p\|}.$$

Pick a small angle $0 < \alpha_{\min} < \pi/2$, and let

$$\mathcal{B}_{\alpha_{\min}}(t_j) = \{q \in \mathcal{B}(t_j) : \sin(\alpha_q) \geq \sin(\alpha_{\min})\}.$$

(In other words, $\mathcal{B}_{\alpha_{\min}}(t)$ contains all points in \mathcal{B} that are not in a conical region of half-width α_{\min} , as shown in Figure 4). For all $q \in \mathcal{B}(t_j)$, $\cos(\alpha'_q) \leq \cos(\alpha_q)$; moreover, for all $q \in \mathcal{B}_{\alpha_{\min}}$,

$$\begin{aligned} \cos(\alpha'_q) &\leq \cos(\alpha_q) - \sin(\alpha_{\min})(\alpha'_q - \alpha_q) \leq \\ &\leq \cos(\alpha_q) - \frac{\delta \sin(\alpha_{\min})^2}{\|q-p\|} \leq \cos(\alpha_q) - \frac{\delta \sin(\alpha_{\min})^2}{\text{diam}(\mathcal{Q}) + \delta}. \end{aligned}$$

Hence, summing over all $q \in \mathcal{B}(t_j)$, we get:

$$\sum_{q \in \mathcal{B}(t_j)} \cos(\alpha'_q) \leq \sum_{q \in \mathcal{B}(t_j)} \cos(\alpha_q) - \sum_{q \in \mathcal{B}_{\alpha_{\min}}(t_j)} \frac{\delta \sin(\alpha_{\min})^2}{\text{diam}(\mathcal{Q}) + \delta} \quad (17)$$

Observe now that in any case

$$\left| \sum_{q \in \mathcal{B}(t_j)} \cos(\alpha_q) \right| \leq 1,$$

(it is zero in case $p^*(t_j) \notin \mathcal{B}(t_j)$, and bounded in absolute value by one if $p^*(t_j) \in \mathcal{B}(t_j)$). Therefore, rearranging equation (17):

$$\text{card}(\mathcal{B}_{\alpha_{\min}}(t_j)) \frac{\delta \sin(\alpha_{\min})^2}{\text{diam}(\mathcal{Q}) + \delta} \leq \sum_{q \in \mathcal{B}(t_j)} \cos(\alpha_q) - \sum_{q \in \mathcal{B}(t_j)} \cos(\alpha'_q) \leq 2$$

Solving this inequality with respect to δ we get:

$$\delta \leq \frac{2 \text{diam}(\mathcal{Q})}{\text{card}(\mathcal{B}_{\alpha_{\min}}(t_j)) \sin(\alpha_{\min})^2 - 2}. \quad (18)$$

Since (i) α_{\min} is a positive constant, (ii) $p^*(t_j)$ is in the interior of \mathcal{Q} , and (iii) $\lim_{j \rightarrow \infty} \text{card}(\mathcal{B}(t_j)) = +\infty$, the right hand side of (18) converges to zero with probability one. Since the bound holds for all points $p \in C(t_j)$, for all $j \in \mathbb{N}$, the claim follows. \blacksquare

In the previous proposition, we have proven that $\|p^*(t_{j+1}) - p^*(t_j)\|$ tends to zero a.s. as $j \rightarrow \infty$, under some natural assumptions on the distribution φ and its support. Unfortunately this is not sufficient to prove that the sequence $\{p^*(t_j)\}_{t_j \in \mathbb{N}}$ is Cauchy; convergence of the sequence is however ensured by the following

Proposition 9 *The sequence $\{p^*(t_j)\}_{j \in \mathbb{N}}$ converges almost surely.*

Proof: Since the sequence $\{p^*(t_j)\}_{j \in \mathbb{N}}$ takes value in a compact set (the closure of $\text{Supp}(\varphi)$), by the Bolzano-Weirstrass theorem there exists a subsequence converging to a limit point \hat{p}^* in the compact set. Construct from $\{p^*(t_j)\}_{j \in \mathbb{N}}$ a maximal subsequence converging to \hat{p}^* , and call J the set of indices of this maximal subsequence. If the original sequence $\{p^*(t_j)\}_{j \in \mathbb{N}}$ is not converging to \hat{p}^* , then there exists an $L > 0$ such that $\|p^*(t_j) - \hat{p}^*\| \geq L$, for any $j \in \mathbb{N} \setminus J$, and this set of indices is unbounded. Take $\epsilon := L/3 > 0$. We have that $\|p^*(t_{k-1}) - \hat{p}^*\| \leq \epsilon$ for any sufficiently large $(k-1) \in J$; moreover, $\|p^*(t_{k-1}) - p^*(t_k)\| \leq \epsilon$, a.s. by Proposition (8). Choose a sufficiently large $(k-1) \in J$, such that $k \in \mathbb{N} - J$ (this is always possible since the complementary set of J is unbounded by the assumption of non-convergence). But now, we have

$$L \leq \|p^*(t_k) - \hat{p}^*\| \leq \|p^*(t_k) - p^*(t_{k-1})\| + \|p^*(t_{k-1}) - \hat{p}^*\| \leq \frac{2}{3}L,$$

which is a contradiction. \blacksquare

4.3 Convergence to the generalized median

By the discussion in the previous section, we know that the reference points of all agents that visit an unbounded number of targets converge to a well-defined limit. So do, trivially, the reference points of all agents that visit a bounded set of targets. Hence, we know that the sequence of reference points $p_i^*(t_j)$ converges to a limit \hat{p}_i^* , almost surely for all $i \in \{1, \dots, m\}$. Let us denote by $\mathcal{V}_i(p^*(t_j))$ the Voronoi region associated to the generator $p_i^*(t_j)$, and by $\mathcal{V}_i(\hat{p}^*)$ the Voronoi region corresponding to the limit point \hat{p}_i^* .

Proposition 10 *If the limit reference points $\hat{p}^* = (\hat{p}_1^*, \dots, \hat{p}_2^*)$ are distinct, then the sequence of Voronoi partitions $\{\mathcal{V}(p^*(t_j))\}_{j \in \mathbb{N}}$ converges to the Voronoi partition generated by the limit of reference points, i.e.,*

$$\lim_{j \rightarrow \infty} \mathcal{V}_i(p^*(t_j)) = \mathcal{V}_i(\hat{p}^*), \quad \text{a.s.}$$

Proof: The boundaries of regions in a Voronoi partition are algebraic curves that depend continuously on the generators, as long as these are distinct. Hence, under this assumption, almost sure convergence of the generators implies the almost sure convergence of the Voronoi regions. ■

As a next step, we wish to understand what is the relation between the asymptotic reference positions and their associated Voronoi regions. More precisely, let $\mathcal{A} \subset \{1, \dots, m\}$ be the subset of indices of agents that visit an unbounded number of targets; we want to prove that \hat{p}_i^* is indeed the generalized median \bar{p}_i associated to agent i , with respect to the limiting set $\mathcal{V}_j(\hat{p}^*)$ and distribution $\varphi(x)$, $\forall i \in \mathcal{A}$. First we need the following technical result.

Lemma 11 *Let $\{f_i\}_{i \in \mathbb{N}} : \mathcal{Q} \rightarrow \mathbb{R}$ be a sequence of strictly convex continuous functions, defined on a common compact subset $\mathcal{Q} \subset \mathbb{R}^n$. Assume that each f_i has a unique $x_i := \arg \min_x f_i$ belonging to the interior of Ω for any i and that this sequence of function converges uniformly to a continuous strictly convex function f admitting a unique minimum point \bar{x} belonging also to the interior of Ω . Then $\lim_{i \rightarrow \infty} x_i = \bar{x}$.*

Proof: Since $\{f_i\}_{i \in \mathbb{N}}$ converges uniformly to f , then for any $\epsilon > 0$, there exists an $I(\epsilon)$ such that for any $i \geq I(\epsilon)$, $\|f_i - f\| \leq \epsilon$ uniformly in $x \in \Omega$. Let $m := f(\bar{x})$, the minimum value achieved by f . Consider the set

$$U_\epsilon := \{x \in \Omega \text{ such that } f(x) \leq m + 2\epsilon\}.$$

Since f is strictly convex, for ϵ sufficiently small, U_ϵ will be a compact subset contained in Ω . Moreover, since f is strictly convex, we have that U_ϵ is strictly included in $U_{\epsilon'}$, whenever $\epsilon' > \epsilon$ and both are sufficiently small. It is also clear that $\lim_{\epsilon \rightarrow 0} U_\epsilon = \bar{x}$ (nested strictly decreasing sequence of compact subsets all containing the point \bar{x}). If $\|f_i - f\| \leq \epsilon$, we claim that $x_i \in U_\epsilon$; we prove this by contradiction. Since $f_i \geq f - \epsilon$, if x_i does not belong to U_ϵ , then $\min(f_i) = f_i(x_i) > m + \epsilon$ (this is just because U_ϵ is simply the set where the function $(f - \epsilon) \leq m + \epsilon$). But since $f_i \leq f + \epsilon$ it turns out that $f_i(\bar{x}) \leq m + \epsilon < \min(f_i)$, which is a contradiction. ■

We conclude this section with the following:

Proposition 12 *Assume that all agents in \mathcal{A} are able to return infinitely often to their reference point between visiting two targets. Then, the limit reference points of such agents coincide, almost surely, with the generalized medians of their limit Voronoi regions, i.e.,*

$$\hat{p}_i^* = \arg \min_{p \in \Omega} \int_{\mathcal{V}_i(\hat{p}^*)} \varphi(q) dq, \quad \text{a.s.,} \quad \forall i \in \mathcal{A}.$$

Proof: For any $i \in \mathcal{A}$, define the functions $f_{i,t_j}(p) := \frac{1}{j} \sum_{q \in \mathcal{B}_i(t_j)} \|p - q\|$ and $f_i(p) := \int_{\mathcal{V}_i(\hat{p}^*)} \|p - q\| \varphi(q) dq$. These functions are continuous and well defined over Ω . We restrict their domains of definition to the compact set $\mathcal{Q} = \text{Supp}(\varphi)$. These functions are also strictly convex and have unique minima in the interior of \mathcal{Q} . Let us notice that, with our previous notation, we have that $p_i^*(t_j) = \arg \min f_{i,t_j}(p)$ and $\bar{p}_i = \arg \min f_i(p)$. Observe that the functions $f_{i,t_j}(p)$ and $f_i(p)$ can be considered random variables with respect to a probability space whose space of events coincide with all possible realizations of target sequences. Consider a restriction of these random variable to a new probability space whose space of events coincide with all possible realizations of target sequences, for which the corresponding FT points converge to a limiting point. On this new probability space the random variable $f_i(p)$ becomes a deterministic function which is the expected value of the random variables $f_{i,t_j}(p)$. Since \mathcal{Q} is compact, it is immediate to see that $f_{i,t_j}(p)$ have finite expectation and variance over this reduced probability space, and by the Strong Law of Large Numbers we can conclude that almost surely (over this reduced probability space) $f_{i,t_j}(p)$ converge pointwise to $f_i(p)$. To show that $f_{i,t_j}(p)$ converges pointwise to $f_i(p)$ over the original probability space, it is sufficient to observe the following. The original probability space is the probability space whose space of events coincide with all possible realizations of target sequences. We already know that almost surely the FT points associated to any possible realization of target sequences will converge. So we can fiber the space of events of the first probability space into spaces of events of reduced probability spaces, except for a set of measure zero. This is sufficient to prove that $f_{i,t_j}(p)$ converge pointwise to $f_i(p)$ almost surely with respect to all possible realizations of target sequences.

Now that we have proved that almost surely the sequence $\{f_{i,t_j}(p)\}_{t_j \in \mathbb{R}}$ converges pointwise to $f_i(p)$, we prove that it does converge uniformly. To do this, we use a theorem, usually attributed to Dini-Arzelà' which state the following: an equicontinuous sequence of functions converges uniformly to a continuous function on a compact set \mathcal{Q} if and only if it converges pointwise to a continuous function on the same compact set. Our sequence $f_{i,t_j}(p)$ is equicontinuous if $\forall \epsilon > 0$ and $\forall p \in \mathcal{Q}$ there exists a $\delta > 0$ such that for all $j \in \{1, \dots, n\}$ and for all $p' \in \mathcal{Q}$ with $\|p' - p\| \leq \delta$, we have $\|f_{i,t_j}(p) - f_{i,t_j}(p')\| \leq \epsilon$; observe that δ is independent on j , while in general it will depend on ϵ and on p . Now we have

$$\|f_{i,t_j}(p) - f_{i,t_j}(p')\| \leq \frac{1}{j} \sum_{q \in \mathcal{B}_i(t_j)} \left| \|p - q\| - \|p' - q\| \right|.$$

Using

$$\|q - p'\| = \|q - p + p - p'\| \leq \|q - p\| + \|p - p'\|$$

and

$$\|q - p\| = \|q - p' + p' - p\| \leq \|q - p'\| + \|p - p'\|,$$

it is immediate to see that

$$\|f_{i,t_j}(p) - f_{i,t_j}(p')\| \leq \frac{1}{j} \sum_{q \in \mathcal{B}_i(t_j)} \|p - p'\| \leq \|p - p'\|.$$

So it is sufficient to take $\delta = \epsilon$ in the previous definition and δ does not depend on j . So the sequence is equicontinuous and the pointwise convergence is upgraded to uniform convergence.

We already know that almost surely the points $p_i^*(t_j)$ do converge to points \hat{p}_i^* (Proposition (9)); therefore we can claim that $\hat{p}_i^* = \bar{p}_i$, simply applying Lemma (11), which requires the uniform convergence. Thus we can claim that the reference position of each agent which services infinitely many targets following our algorithm converges to the generalized median of its Voronoi region, almost surely. ■

4.4 Fairness

In this section, we prove that, as long as φ is strictly positive over a convex set, both policies introduced in Section 3 are fair.

Proposition 13 (Fairness) *If $\mathcal{Q} = \text{Supp}(\varphi)$ is convex, all agents eventually visit an unbounded number of targets, almost surely, i.e.,*

$$\lim_{t \rightarrow +\infty} \text{card}(\mathcal{B}_i(t)) = +\infty, \quad \text{a.s.}, \quad \forall i \in \{1, \dots, m\}.$$

Proof: Under either policy, each agent will pursue and is guaranteed to service the nearest target in its own Voronoi region. Hence, in order to show that an agent services an unbounded number of targets, it is sufficient to show that the probability that the next target be generated within its Voronoi region remains strictly positive, i.e., that

$$\int_{\mathcal{V}_i(p(t))} \varphi(q) dq > 0.$$

Since \mathcal{Q} is convex, and all agents move towards the nearest target, at least initially, all agents will eventually enter \mathcal{Q} , and remain within it. Let us denote with P_{ij} the probability that agent i visits the j -th target. For any i this probability is always strictly positive. Indeed, even if it happens that some agents are servicing simultaneously the same target (simply because the service request appears at the boundary of two different Voronoi regions), this does not mean that their reference points have to coincide. For the reference points of two agents eventually to coincide, it must happen that they are servicing infinitely many often and simultaneously the same target. Since the boundaries of Voronoi regions have measure zero and φ is a continuous distribution without singular components, we can claim that $\lim_{j \rightarrow \infty} P_{ij} > 0$ almost surely. Now call $P = \min_{i=1, \dots, m} \lim_{j \rightarrow \infty} P_{ij}$. Then P is strictly positive almost surely. Therefore, the probability that the i -th agent *does not* visit an unbounded number of targets is bounded from above by $\lim_{j \rightarrow \infty} \prod_{k=1}^j (1 - P)^k = 0$, almost surely. ■

4.5 Efficiency

In this section, we will prove that the system time provided by either one of the algorithms in Section 3 converges to a critical point (either a saddle point or a local minimum) with high probability as $\lambda \rightarrow 0$.

In the preceding sections, we have proved that—as long as each agent is able to return to its reference point between servicing two targets, infinitely often—the reference points $p_i^*(t_j)$ converge to points \hat{p}_i^* , which generate a MVT. In such case, we know that the average time of service will converge to

$$\bar{T}_\pi = \int_{\Omega} \min_{i=1,\dots,m} \|\hat{p}_i^* - q\| \varphi(q) dq = \sum_{i=1}^m \int_{\mathcal{V}_i(\hat{p}^*)} \|\hat{p}_i^* - q\| \varphi(q) dq. \quad (19)$$

Consider now functions \mathcal{H}_m of the form:

$$\mathcal{H}_m(p_1, \dots, p_m) = \int_{\Omega} \min_{i=1,\dots,m} \|p_i - q\| \varphi(q) dq = \sum_{i=1}^m \int_{\mathcal{V}_i(p)} \|p_i - q\| \varphi(q) dq. \quad (20)$$

Observe that \bar{T}_π belongs to the class of functions of the form \mathcal{H}_m where each point p_i is constrained to be the generalized median of the corresponding Voronoi region (i.e., $\mathcal{V}(p)$ is a MVT).

We want to prove that \bar{T}_π is a critical point of \mathcal{H}_m . To do so, we consider an extension of \mathcal{H}_m , i.e. a functional \mathcal{K}_m defined as follows:

$$\mathcal{K}_m(p_1, \dots, p_m, \mathcal{V}_1, \dots, \mathcal{V}_m) := \sum_{i=1}^m \int_{y \in \mathcal{V}_i} \|y - p_i\| \varphi(y) dy.$$

Observe that in this case the regions $\{\mathcal{V}_i\}_{i=1,\dots,m}$ are not restricted to form a MVT with respect to the generators $\{x_i\}_{i=1,\dots,m}$. Thus we can view the functional \mathcal{H}_m we are interested in as a constrained form of the unconstrained functional \mathcal{K}_m . It turns out therefore that critical points of \mathcal{K}_m are also critical points of \mathcal{H}_m . With respect to critical points of \mathcal{K}_m we have the following result:

Proposition 14 *Let $\{p_i\}_{i=1,\dots,m}$ denote any set of m points belonging to $\text{Supp}(\varphi)$ and let $\{\mathcal{V}_i\}_{i=1,\dots,m}$ denote any tessellation of $\text{Supp}(\varphi)$ into m regions. Moreover, let us define \mathcal{K}_m as above.*

Then a sufficient condition for $\{p_1, \dots, p_m, \mathcal{V}_1, \dots, \mathcal{V}_m\}$ to be a critical point (either a saddle point or a local minimum), is that the \mathcal{V}_i 's are the Voronoi regions corresponding to the p_i 's, and, simultaneously, the p_i 's are the generalized median of the corresponding \mathcal{V}_i 's.

Proof: Consider first the variation of \mathcal{K}_m with respect to a single point, say p_i . Now let v be a vector in \mathbb{R}^2 , such that $p_i + \epsilon v \in \Omega$. Then we have

$$\mathcal{K}_m(p_i + \epsilon v) - \mathcal{K}_m(p_i) = \int_{y \in \mathcal{V}_i} \{\|y - p_i - \epsilon v\| - \|y - p_i\|\} \varphi(y) dy,$$

where we have not listed the other variables on which \mathcal{K}_m depends since they remain constant in this variation. By the very form of this variation, it is clear that if the point p_i is the generalized median for the *fixed* region \mathcal{V}_i , we will have that $\mathcal{K}_m(p_i + \epsilon v) - \mathcal{K}_m(p_i) > 0$, for any v . Now consider the points $\{p_i\}_{i=1,\dots,m}$ fixed and consider a tessellation $\{\mathcal{U}_i\}_{i=1,\dots,m}$ different from the Voronoi regions $\{\mathcal{V}_i\}_{i=1,\dots,m}$ generated by the points p_i 's. We compare the value of $\mathcal{K}_m(p_1, \dots, p_m, \mathcal{V}_1, \dots, \mathcal{V}_m)$, with the value of $\mathcal{K}_m(p_1, \dots, p_m, \mathcal{U}_1, \dots, \mathcal{U}_m)$. Consider those y which belong to the Voronoi region \mathcal{V}_j generated by p_j , and possibly not to the Voronoi region of another p_i . Anyway, since \mathcal{U}_i is not a Voronoi tessellation, it can happen that in any case these y belong to \mathcal{U}_i . Thus for these particular y 's we have $\varphi(y)\|y - p_j\| \leq \varphi(y)\|y - p_i\|$. Moreover, since $\{\mathcal{U}_i\}_{i=1,\dots,m}$ are not the Voronoi tessellation associated to the p_i 's, the last inequality must be strict over some set of positive measure. Thus we have that $\mathcal{K}_m(p_1, \dots, p_m, \mathcal{V}_1, \dots, \mathcal{V}_m) < \mathcal{K}_m(p_1, \dots, p_m, \mathcal{U}_1, \dots, \mathcal{U}_m)$, and therefore \mathcal{K}_m is minimized, keeping fixed the p_i 's exactly when the subset \mathcal{V}_i 's are chosen to be the Voronoi regions associated with the point p_i 's. ■

By the previous proposition and by the fact that critical points of the unconstrained functional \mathcal{K}_m are also critical points of the constrained functional \mathcal{H}_m , we have that the MVT are always critical points for the functional \mathcal{H}_m , and in particular \bar{T} is either a saddle point or a local minimum for the functional \mathcal{H}_m .

Before we conclude, we need one last intermediate result.

Proposition 15 *Each agent will be able to return to its reference point before the generation of a new service request infinitely often with high probability as $\lambda \rightarrow 0$.*

Proof: Let t_1 be such that $\mathcal{B}_i(t_1) \neq \emptyset$, for all $i \in \{1, \dots, m\}$. Such time exists, almost surely, because of the fairness of the proposed policies. At time t_1 all agents will be within \mathcal{Q} . Let $n_1 = \text{card}(\mathcal{D}(t_1))$ be the total number of outstanding targets at time t_1 . An upper bound on the time needed to visit all targets in $\mathcal{D}(t_1)$ is $n_1(\text{diam}(\mathcal{Q}))$. When there are no outstanding targets, agents move to their reference points, reaching them in at most $\text{diam}(\mathcal{Q})$ units of time.

The time needed to service the initial targets and go to the reference configuration is hence bounded by $t_{\text{ini}} \leq t_1 + (n_1 + 1)\text{diam}(\mathcal{Q})$. The probability that at the end of this initial phase the number of targets is reduced to zero is

$$P[n(t_{\text{ini}}) = 0] = \exp(-\lambda(t_{\text{ini}} - t_1)) \geq \exp(-\lambda(n_0 + 1)\text{diam}(\mathcal{Q})),$$

that is, $P[n(t_{\text{ini}}) = 0] \rightarrow 1^-$ as $\lambda \rightarrow 0^+$. As a consequence, after an initial transient, all targets will be generated with all agents waiting at their reference points, and an empty demand queue, with high probability as $\lambda \rightarrow 0^+$. ■

We can now conclude with following:

Theorem 16 (Efficiency) *The system time provided by the no-communication policy π_{nc} and by the sensor-based policy π_{sb} converges to a critical point (either a saddle point or a local minimum) with high probability as $\lambda \rightarrow 0$.*

Proof: Combining results in Propositions 9 and 12 we conclude that the reference points of all agents that visit an unbounded number of targets converge to a MVT, almost surely—provided agents can return to the reference point between visiting targets. Moreover, the fairness result in Proposition 13 shows that in fact all agents do visit an unbounded number of targets almost surely; as a consequence, Proposition 14 the limit configuration is indeed a critical point for the system time. Since agents return infinitely often to their reference positions with high probability as $\lambda \rightarrow 0$, the claim is proven. ■

Thus we have proved that the suggested algorithm enables the agents to realize a coordinated task, such that “minimizing” the cost function without explicit communication, or with mutual position knowledge only. Let us underline that, in general, the achieved critical point strictly depends on the initial positions of the agents inside the environment Ω . It is known that the function \mathcal{H}_m admits (not unique, in general) *global minima*, but the problem to find them is *NP-hard*.

Remark 17 *We can not exclude that the algorithm so designed will converge indeed to a saddle point instead of a local minimum. This is due to the fact that the algorithm provides a sort of implementation of the steepest descent method, where, unfortunately we are not following the steepest direction of the gradient of the function \mathcal{H}_m , but just the gradient with respect to one of the variables. For a broader point of view of steepest descent in this framework see for instance [26].*

On the other hand, since the algorithm is based on a sequence of targets and at each phase we are trying to minimize a different cost function, it can be proved that the critical points reached by this algorithm are *no worse* than the critical points reached knowing a priori the distribution φ . This is a remarkable result proved in a different context in [27], where it is also presented an example in which the use of a sample sequence provides a better result (with probability one) than the a priori knowledge of φ . In that specific example the algorithm with the sample sequence does converge to a global minimum, while the algorithm based on the a priori knowledge of the distribution φ gets stuck in a saddle point.

4.6 A comparison with algorithms for vector quantization and centroidal Voronoi tessellations

The use of Voronoi tessellations is ubiquitous in many fields of science, ranging from operative research, animal ethology (territorial behaviour of animals), computer science (design of algorithms), to numerical analysis (construction of adaptive grids for PDEs and general quadrature rules), and algebraic geometry (moduli spaces of abelian varieties). For a detailed account of possible applications see for instance the book [26]. In the available literature, most of the analysis is devoted to applications of centroidal Voronoi tessellations, i.e., Voronoi tessellation such that

$$p_i = \arg \min_{s \in \mathbb{R}^n} \int_{\mathcal{V}_i(p_1, p_2, \dots, p_m)} \|s - q\|^2 \varphi(q) dq, \quad \forall i \in \{1, \dots, m\}.$$

A popular algorithm due to Lloyd [18] is based on the iterative computation of centroidal Voronoi tessellations. The algorithm can be summarized as follows. Pick m generator points, and consider a large number n of samples from a certain distribution. At each step of the algorithm generators are moved towards the centroid of the samples inside their respective Voronoi region. The algorithm terminates when each generator is within a given tolerance from the centroid of samples in its region, thus obtaining a centroidal Voronoi tessellation weighted by the sample distribution. There is also a continuous version of the algorithm, which requires the a priori knowledge of a spatial density function, and computation of the gradient of the polar moments of the Voronoi regions with respect to the positions of the generators. An application to coverage problems in robotics and sensor networks of Lloyd's algorithm is available in [13].

The algorithms we introduced in this paper are more closely related to an algorithm due to MacQueen [19], originally designed as a simple online adaptive algorithm to solve clustering problems, and later used as the method of choice in several vector quantization problems where little information about the underlying geometric structure is available. MacQueen's algorithm can be summarized as follows. Pick m generator points. Then iteratively sample points according to the probability density function φ . Assign the sampled point to the nearest generator, and update the latter by moving it in the direction of the sample. In other words, let q_j be the j -th sample, let $i^*(j)$ be the index of the nearest generator, and let $c = (c_1, \dots, c_m)$ be a counter vector, initialized to a vector of ones. The update rule takes the form

$$p_{i^*(j)} \leftarrow \frac{c_{i^*(j)} q_j + p_{i^*(j)}}{c_{i^*(j)} + 1},$$

$$c_{i^*(j)} \leftarrow c_{i^*(j)} + 1.$$

The process is iterated until some termination criterion is reached. Compared to Lloyd's algorithm, MacQueen's algorithm has the advantage to be a learning adaptive algorithm, not requiring the a priori knowledge of the distribution of the objects, but rather allowing the online generation of samples. It can be recognized that the update rule in MacQueen's algorithm corresponds to moving the generator points to the centroids of the samples assigned to them.

The algorithm we propose is very similar in spirit to MacQueen's algorithm, however, there is a remarkable difference. MacQueen's algorithm deals with centroidal Voronoi tessellations, thus with the computation of k -means. Our algorithm instead is based on MVT, and on the computation of k -medians. In general, very little is known for Voronoi diagrams generated using simply the Euclidean distance instead of the square of the Euclidean distance. For instance, the medians of a sequence of points can exhibit a quite peculiar behavior if compared to the one of the means. Consider the following example. Given a sequence of points $\{q_i\}_{i \in \mathbb{N}}$ in a compact set $K \subset \mathbb{R}^2$, we can construct the

induced sequence of means:

$$m_N := \frac{1}{N} \sum_{i=1}^N q_i$$

and analogously the induced sequence of FT points we considered in the previous sections. Call FT_N the FT point corresponding to the first N points of the sequence $\{q_i\}_{i \in \mathbb{N}}$. We want to point out that induced sequence $\{m_j\}$ and $\{FT_j\}$ have a very different behaviour. Indeed, the induced sequence of means will always converge as long as the points q_j s belong to a compact set. To see this, just observe that if $\text{diam}(K) \leq L$, then $\|m_j\| \leq L$. Moreover, it is immediate to see that $\|m_{N+1} - m_N\| \leq \frac{2L}{N}$. Then one can conclude using the same argument of Theorem (9). On the other hand, one can construct a special sequence of points q_j s in a compact set K for which the induced sequence of FT points does not converge. This is essentially due to the fact that while the contribution of each single point q_j in moving the position of the mean decreases as j increases, this could not happen in the case of the median. To give a simple example, start with the following configuration of points in \mathbb{R}^2 : $q_1 = (1, 0)$, $q_2 = (-1, 0)$, $q_3 = (0, 1)$ and $q_4 = (0, -1)$. Then the sequence of points q_j s continue in the following way: $q_k = (0, 1)$ if $k > 4$, and k odd, $q_k = (0, -1)$ if $k > 4$ and k is even. Using the characterization of FT points, it is immediate to see that $FT_k = (0, 0)$ for $k > 4$ and k even, while $FT_k = (0, \tan(\pi/6))$ for $k > 4$ and k odd, so the induced sequence can not converge. This phenomenon can not happen to the sequence of means, which is instead always convergent. Therefore, it should be clear that the use of MVT instead of centroidal Voronoi tessellations makes much more difficult to deal with the technical aspects of the algorithm such as its convergence.

5 A game-theoretic point of view

In this section we provide an analysis of the proposed algorithm from the point of view of game theory. In particular, we frame our presentation on the works [7], [28], in which the point of view of game theory has been introduced in the study of cooperative control and strategic coordination of decentralized networks of multi-agents systems. In this section we prove that our algorithm provides a pure Nash equilibrium in a multi-player game where each agent is interested in maximizing its own utility. On the other hand, our multi-player game formulation is much simpler than the usual framework considered in the literature, since there will be no negotiation mechanism among the agents. Despite this fact it turns out that in this example, just trying to maximize their own utility function the agents will indeed maximize a *different global utility function*.

In this section we view the agents as *rational autonomous decision makers* trying to maximize their own utility function. The utility function of agent i , denoted by U_i is simply the *expected* number of service requests handled by agent

i within a certain time horizon, bounded or unbounded, where $i = 1, \dots, m$. We assume that the stochastic process for generating targets is the one already described in the previous sections. It is obvious that any of the utility function U_i is a function of the policy vector $\pi := \{\pi_1, \dots, \pi_m\}$ consisting of all the policies followed by each agent. In general the space of all policies Π is just an uncountable set containing all conceivable policies chosen by an agent and it does not have any other structure. Thus we assume that the policy space Π_i of agent i is equal to a fixed policy space Π which is independent on i , this for any $i = 1, \dots, m$. Therefore the policy vector $\pi \in \Pi^m$. Denoting with π_{-j} : the policy specification of all the agents, except agent j , i.e. $\pi_{-j} := (\pi_1, \dots, \pi_{j-1}, \pi_{j+1}, \dots, \pi_m)$ we may write policy vector π as (π_j, π_{-j}) . Using this notation we can formulate the following definition adapted by [7]:

Definition 18 *A policy vector π^* is called a pure Nash equilibrium if for all $j = 1, \dots, m$:*

$$U_j(\pi_j^*, \pi_{-j}^*) = \max_{\pi_j \in \Pi} U_j(\pi_j, \pi_{-j}^*). \quad (21)$$

Moreover, a policy vector π is called efficient if there is no other policy vector that yields higher utilities to all agents.

Under the target generation assumptions followed so far, we have the following:

Proposition 19 *Let us call $\tilde{\pi}_j$ the policy assignment for agent j corresponding to our algorithm. Then the policy vector $\tilde{\pi} := \{\tilde{\pi}_j\}_{j=1, \dots, m}$ is an efficient pure Nash equilibrium for the given agents utilities.*

Proof: It is immediate to see that policy vector $\tilde{\pi}$ satisfies equation (21), and thus is a pure Nash equilibrium. It is also clear that cannot be any other policy assignment which yields *strictly higher* utilities to all agents, simply because when there is a new outstanding service request all agents move directly toward that location, trying to satisfy it as if there were no other agent in the environment. ■

On the other hand, observe that we do *not* claim in the previous proposition that our algorithm provides the unique efficient pure Nash equilibrium for the given set of utilities function. For instance to find other efficient pure Nash equilibria it is sufficient to modify the algorithm during the initial phases, when the FT points are not uniquely determined. These modifications produce different policy vectors which are anyway all efficient pure Nash equilibria, as it is immediate to see.

Our game-theoretic formulation of the given algorithm belongs to a class of multi-player games called potential games. In a potential game the difference in the utility reached by any of the agents for two different policy choices, when the policies of the other agents are kept fixed, can be measured by a potential function that depends on the policy vector and not on the label of any agent. The fact that our game-theoretic formulation belongs to this class is obvious since the functional form of the utility function is the same for each agent, that

is $U_j = U$ for any agent j . So in this case, as a potential function we take U . The formal definition is as follows:

Definition 20 A potential game is the set $\{U_1, \dots, U_m, \psi\}$ consisting of agent utilities $U_1(\pi), \dots, U_m(\pi)$ and a potential function $\psi : \Pi^m \rightarrow \mathbb{R}$, such that for every agent a_j and for any policy assignments $\pi'_j, \pi''_j \in \Pi_j$ and $\pi_{-j} \in \prod_{k \neq j} \Pi_k$:

$$U_j(\pi'_j, \pi_{-j}) - U_j(\pi''_j, \pi_{-j}) = \psi(\pi'_j, \pi_{-j}) - \psi(\pi''_j, \pi_{-j}).$$

An extension of this concept is provided by:

Definition 21 An ordinal potential game is the set $\{U_1, \dots, U_m, \psi\}$ consisting of agent utilities $U_1(\pi), \dots, U_m(\pi)$ and a potential function $\psi : \Pi^m \rightarrow \mathbb{R}$, such that for every agent a_j and for any policy assignments $\pi'_j, \pi''_j \in \Pi_j$ and $\pi_{-j} \in \prod_{k \neq j} \Pi_k$:

$$U_j(\pi'_j, \pi_{-j}) - U_j(\pi''_j, \pi_{-j}) > 0 \text{ if and only if } \psi(\pi'_j, \pi_{-j}) - \psi(\pi''_j, \pi_{-j}) > 0.$$

At this point, we are ready to introduce the global utility function for this multi-player game formulation. The global utility function for this game is given by $U_g(\pi) = -\overline{T}_\pi$, where \overline{T}_π is the *system time* under policy vector π . An important aspect in the game-theoretic approach is to understand to which extent the utility functions of the individual players are compatible with the global utility function. To this aim the following definition has been introduced in [7]:

Definition 22 The set of agents utilities $\{U_j(\pi)\}_{j=1, \dots, m}$ is aligned with the global utility $U_g(\pi)$ iff the set $\{U_1, \dots, U_m, U_g\}$ forms an ordinal potential game, with potential function given by U_g .

It is immediate to prove the following:

Proposition 23 The given agents utilities are aligned with the global utility function, in this game-theoretic formulation of the algorithm.

Proof: Let us focus on one agent, say agent j . If its utility function increases, it is able to service a bigger number of targets in the given time horizon. It can happen that some other agent will have a corresponding decrease in their utility function, and this happens exactly when the agent j due to a policy change will service some targets more rapidly. This in turn, will increase anyway U_g . ■

In general, it is true that alignment does not prevent pure Nash equilibria from being suboptimal from the point of view of the global utility. Moreover, even *efficient* pure Nash equilibria (i.e. pure Nash equilibria which yield the highest utility to all agents) can be suboptimal from the perspective of the global utility function. Such a phenomenon is indeed what happens in our construction.

Proposition 24 *The policy vector $\tilde{\pi}$ corresponding to the policies realized by our algorithm is an efficient pure Nash equilibrium which is possibly suboptimal from the point of view of global utility.*

Proof: We already know that $\tilde{\pi}$ is an efficient pure Nash equilibrium and from the analysis developed in Section (4) we know that it yields a *critical point* for the system time \overline{T}_π . Thus it corresponds to a *critical point* for $U_g(\pi)$, either a local maximum or a saddle point. On the other hand, as noted for \overline{T}_π , U_g may have in general several local maxima and our algorithm is not guaranteed to converge to a global maximum. ■

6 Numerical results

In this section, we present simulation results showing the performance of the proposed policies for various scenarios.

6.1 Uniform distribution, light load

In the numerical experiments, we first consider $m = 9$, choose \mathcal{Q} as a unit square, and set $\varphi = 1$ (i.e., we consider a spatially uniform target-generation process). This choice allows us to determine easily the optimal placement of reference points, at the centers of a tessellation of \mathcal{Q} into nine equal squares, and compute analytically the optimal system time. In fact, it is known that the expected distance of a point q randomly sampled from a uniform distribution within a square of side L from the center of the square c is

$$\mathbb{E}[\|q - c\|] = \frac{\sqrt{2} + \log(1 + \sqrt{2})}{6}L \approx 0.3826L.$$

The results for a small value of λ , i.e., $\lambda = 0.5$, are presented in Figure 5. The average service time converges to a value that is very close to the theoretical limit computed above, taking $L = 1/\sqrt{m} = 1/3$. In both cases, the reference points converge—albeit very slowly—to the generators of a MVT, while the average system time quickly approaches the optimal value.

6.2 Non-uniform distribution, light load

We also present in Figure 6 results of similar numerical experiments with a non-uniform distribution, namely an isotropic normal distribution centered at $(0.25, 0.25)$, with standard deviation equal to 0.25.

6.3 Uniform distribution, dependency on the target generation rate

An interesting set of numerical experiments evaluates the performance of the proposed policies over a large range of values of the target generation rate λ .

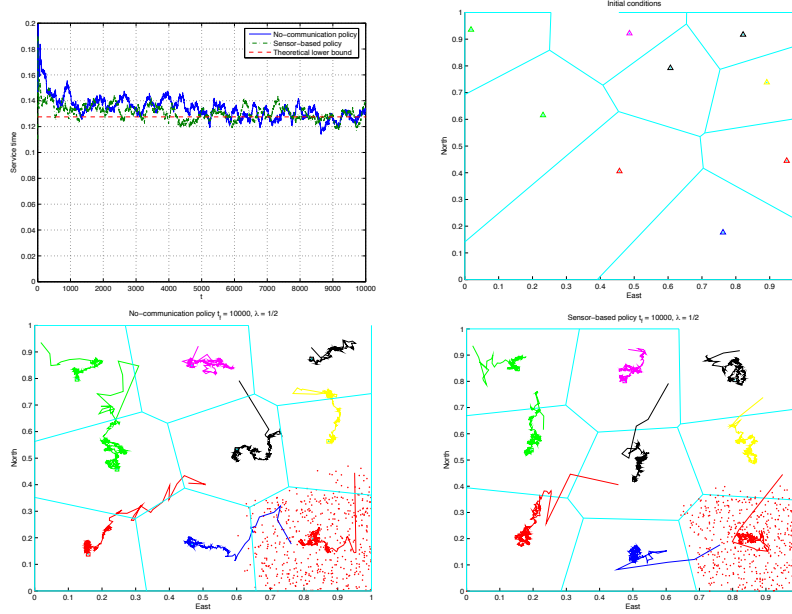


Figure 5: Numerical simulation in the light-load case, for a uniform spatial distribution. Top left: the actual service times as a function of time, for the two policies, compared with the optimal system time. Top right: the initial configuration of the nine agents. Bottom left and right: paths followed by the reference points up to $t = 10^4$ (corresponding to approximately 5,000 targets), using the two policies. The locations of all targets visited by one of the agents are also shown.

In Section 4, we proved the convergence of the system’s behavior to an efficient steady state, with high probability as $\lambda \rightarrow 0$, as confirmed by the simulations discussed above. For large values of λ however, the assumption that vehicles are able to return to their reference point breaks down, and the convergence result is no longer valid. In figure 7 we report results from numerical experiments on scenarios involving $m = 3$ agents, and values of λ ranging from $1/2$ to 32. In the figure, we also report the known (asymptotic) lower bounds on the system time (with 3 agents), as derived in [10], and the system time obtained with the proposed policies in a single-agent scenario.

The performance of both proposed policies is close to optimal for small λ , as expected. The sensor-based policy behaves well over a large range of target generation rates; in fact, the numerical results suggest that the policy provides a system time that is a constant-factor approximation of the optimum, by a factor of approximately 1.6.

However, as λ increases, the performance of the no-communication policy degrades significantly, almost approaching the performance of a single-vehicle system over an intermediate range of values of λ . Our intuition in this phe-

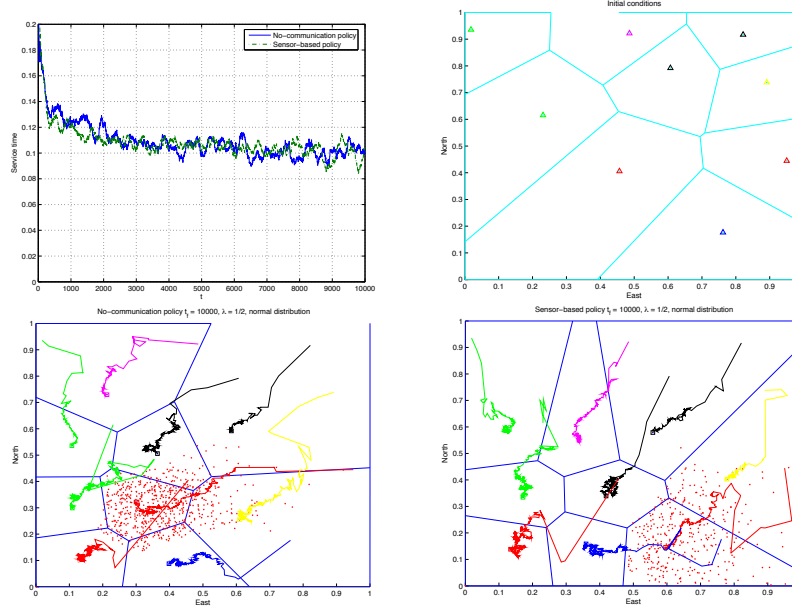


Figure 6: Numerical simulation in the light-load case, for a normal spatial distribution. Top left: the actual service times as a function of time, for the two policies. Top right: the initial configuration of the nine agents. Bottom left and right: paths followed by the reference points up to $t = 10^4$ (corresponding to approximately 5,000 targets), using the two policies. The locations of all targets visited by one of the agents are also shown.

nomenon is the following. As agents do not return to their own reference points between visiting successive targets, their efficiency decreases since they are no longer able to effectively separate regions of responsibility. In practice—unless they communicate and concentrate on their Voronoi region, as in the sensor-based policy—agents are likely to duplicate efforts as they pursue the same target, and effectively behave as a single-vehicle system. Interestingly, this efficiency loss seems to decrease for large λ , and the numerical results suggest that the no-communication policy recovers a similar performance as the sensor-based policy in the heavy load limit. Unfortunately, we are not able at this time to provide a rigorous analysis of the proposed policies for general values of the target generation rate.

7 Conclusions

In this paper we considered two very simple strategies for multiple vehicle routing in the presence of dynamically-generated targets, and analyzed their performance in light load conditions, i.e., when the target generation rate is very




Figure 7: System time provided by the policies proposed in this paper, as a function of the target generation rate λ . The system is composed of three vehicles, and the target points are generated uniformly in the unit square.

small. The strategies we addressed in this paper are based on minimal assumptions on the ability of the agents to exchange information: in one case they do not explicitly communicate at all, and in the other case, agents are only aware of other agents' current location. A possibly unexpected and striking results of our analysis is the following: the collective performance of the agents using such minimal or no-communication strategies is (locally) optimal, and is in fact as good as that achieved by the best known decentralized strategies. Moreover, the proposed strategies do not rely on the knowledge of the target generation process, and makes minimal assumptions on the target spatial distribution; in fact, the convexity and boundedness assumptions on the support of the spatial distribution can be relaxed, as long as path connectivity of the support, and absolute continuity of the distribution are ensured. Also, the distribution needs not be constant: Indeed, the algorithm will provide a good approximation to a local optimum for the cost function as long as the characteristic time it takes for the target generation process to vary significantly is much greater than the relaxation time of the algorithm. In summary, the proposed strategies can be seen as a learning mechanism in which the agents learn the target generation process, and the ensuing target spatial distribution, and adapt their own behavior to it.

The proposed strategies are very simple to implement, as they only require storage of the coordinates of points visited in the past and simple algebraic calculations; the “sensor-based” strategy also require a device to estimate the position of other agents. Simple implementation and the absence of active com-

munication makes the proposed strategies attractive, for example, in embedded systems and stealthy applications. The game-theoretic interpretation of our results also provides some insight into how territorial, globally optimal behavior can arise in a population of selfish but rational individuals even without explicit mechanisms for territory marking and defense.

While we were able to prove that the proposed strategies perform efficiently for small values of the target generation rate, little is known about their performance in other regimes. In particular, we have shown numerical evidence that suggests that the first strategy we introduced, requiring no communication, performs poorly when targets are generated very frequently, whereas the performance of the sensor-based strategy is in fact comparable to that of the best known strategies for the heavy load case.

Extensions of this work will include the analysis and design of efficient strategies for general values of the target generation rate, for different vehicle dynamics models (e.g., including differential constraints on the motion of the agents), and heterogeneous systems in which both service requests and agents can belong to several different classes with different characteristics and abilities.

Acknowledgement The research in this paper was inspired by discussion with Dr. Mahbub Gani, and was performed while the authors were with the Department of Mechanical and Aerospace Engineering at the University of California, Los Angeles.

References

- [1] R. W. Beard, T. W. McLain, M. A. Goodrich, and E. P. Anderson. Coordinated target assignment and intercept for unmanned air vehicles. *IEEE Trans. on Robotics and Automation*, 18(6):911–922, 2002.
- [2] A. Richards, J. Bellingham, M. Tillerson, and J. How. Coordination and control of multiple UAVs. In *Proc. of the AIAA Conf. on Guidance, Navigation, and Control*, Monterey, CA, 2002.
- [3] C. Schumacher, P. R. Chandler, S. J. Rasmussen, and D. Walker. Task allocation for wide area search munitions with variable path length. In *Proc. of the American Control Conference*, pages 3472–3477, Denver, CO, 2003.
- [4] M. Earl and R. D’Andrea. Iterative MILP methods for vehicle control problems. *IEEE Trans. on Robotics*, 21:1158–1167, December 2005.
- [5] W. Li and C. Cassandras. A cooperative receding horizon controller for multivehicle uncertain environments. *IEEE Trans. on Automatic Control*, 51(2):242–257, 2006.
- [6] R.A. Murphey. Target-based weapon target assignment problems. In P.M. Pardalos and L.S. Pitsoulis, editors, *Nonlinear Assignment Problems: Algorithms and Applications*, pages 39–53. Kluwer Academic Publisher, 1999.

- [7] G. Arslan and J.S. Shamma. Autonomous vehicle-target assignment: a game theoretic formulation. Submitted to the IEEE Trans. on Automatic Control, February 2006.
- [8] M. Tanemura and H. Hasegawa. Geometrical models of territory I: Models for synchronous and asynchronous settlement of territories. *Journal of Theoretical Biology*, 82:477–496, 1980.
- [9] H. Psaraftis. Dynamic vehicle routing problems. In B. Golden and A. Assad, editors, *Vehicle Routing: Methods and Studies*, Studies in Management Science and Systems. Elsevier, 1988.
- [10] D. J. Bertsimas and G. J. van Ryzin. A stochastic and dynamic vehicle routing problem in the Euclidean plane. *Operations Research*, 39:601–615, 1991.
- [11] E. Frazzoli and F. Bullo. Decentralized algorithms for vehicle routing in a stochastic time-varying environment. In *Proc. IEEE Conf. on Decision and Control*, Paradise Island, Bahamas, December 2004.
- [12] Y. Liu, J.B. Cruz, and A. G. Sparks. Coordinated networked uninhabited aer vehicles for persistent area denial. In *IEEE Conf. on Decision and Control*, pages 3351–3356, Paradise Island, Bahamas, 2004.
- [13] J. Cortés, S. Martínez, T. Karatas, and F. Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2):243–255, 2004.
- [14] B.J. Moore and K.M. Passino. Distributed balancing of AAVs for uniform surveillance coverage. In *IEEE Conference on Decision and Control*, 2005.
- [15] J. Little. A proof of the queueing formula $l = \lambda w$. *Operations Research*, 9:383–387, 1961.
- [16] D. J. Bertsimas and G. J. van Ryzin. Stochastic and dynamic vehicle routing in the Euclidean plane with multiple capacitated vehicles. *Operations Research*, 41(1):60–76, 1993.
- [17] D. J. Bertsimas and G. J. van Ryzin. Stochastic and dynamic vehicle routing with general interarrival and service time distributions. *Advances in Applied Probability*, 25:947–978, 1993.
- [18] S. P. Lloyd. Least squares quantization in PCM. *IEEE Trans. on Information Theory*, 28(2), 1982.
- [19] J. Mac Queen. Some methods for the classification and analysis of multivariate observations. In L. M. LeCam and J. Neyman, editors, *Proceedings of the Fifth Berkeley Symposium on Math. Stat. and Prob.*, pages 281–297. University of California Press, 1967.

- [20] Z. Drezner, editor. *Facility Location: A Survey of Applications and Methods*. Springer Series in Operations Research. Springer Verlag, New York, 1995.
- [21] P. K. Agarwal and M. Sharir. Efficient algorithms for geometric optimization. *ACM Computing Surveys*, 30(4):412–458, 1998.
- [22] R. Chandrasekaran and A. Tamir. Algebraic optimization: the Fermat-Weber location problem. *Mathematical Programming*, 46:219–224, 1990.
- [23] G. Wesolowsky. The Weber problem: History and perspectives. *Location Science*, 1:5–23, 1993.
- [24] S. P. Fekete, J. S. B. Mitchell, and K. Weinbrecht. On the continuous Weber and k -median problems. In *Proceedings of the Sixteenth Annual Symposium on Computational Geometry (Hong Kong, 2000)*, pages 70–79, New York, 2000. ACM.
- [25] P. Carmi, S. Har-Peled, and M.J. Katz. On the fermat-weber center of a convex object. *Computational Geometry*, 32(3):188–195, 2005.
- [26] A. Okabe, B. Boots, K. Sugihara, and S.N. Chiu. *Spatial tessellations: Concepts and Applications of Voronoi diagrams*. Wiley Series in Probability and Statistics. John Wiley & Sons, Chichester, UK, second edition, 2000.
- [27] M.J. Sabin and R.M. Gray. Global convergence and empirical consistency of the generalized Lloyd algorithm. *IEEE Trans. on Information Theory*, 32(2), March 1986.
- [28] J.S. Shamma and G. Arslan. Dynamic fictitious play, dynamic gradient play, and distributed convergence to Nash equilibria. *IEEE Trans. on Automatic Control*, 50(3):312–327, March 2005.